(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
15 November 2001 (15.11.2001)

PCT

(10) International Publication Number
WO 01/86627 A2

(51) International Patent Classification⁷: G10H

(21) International Application Number: PCT/GB01/01983

(22) International Filing Date: 4 May 2001 (04.05.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

0010967.8	5 May 2000 (05.05.2000)	GB
0010969.4	5 May 2000 (05.05.2000)	GB
0011178.1	9 May 2000 (09.05.2000)	GB
0022164.8	11 September 2000 (11.09.2000)	GB
0030836.1	18 December 2000 (18.12.2000)	GB

(71) Applicant *(for all designated States except US)*: **SSEYO LIMITED** [GB/GB]; Highview House, Charles Square, Bracknell, Berkshire RG12 1DF (GB).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **COLE, John, Tim** [GB/GB]; Sseyo Limited, Highview House, Charles Square, Bracknell, Berkshire RG12 1DF (GB). **COLE, Murray, Peter** [GB/GB]; Sseyo Limited, Highview House, Charles Square, Bracknell, Berkshire RG12 1DF (GB).

(74) Agents: **MAGGS, Michael, Norman** et al.; Kilburn & Strode, 20 Red Lion Street, London WC1R 4PJ (GB).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: AUTOMATED GENERATION OF SOUND SEQUENCES

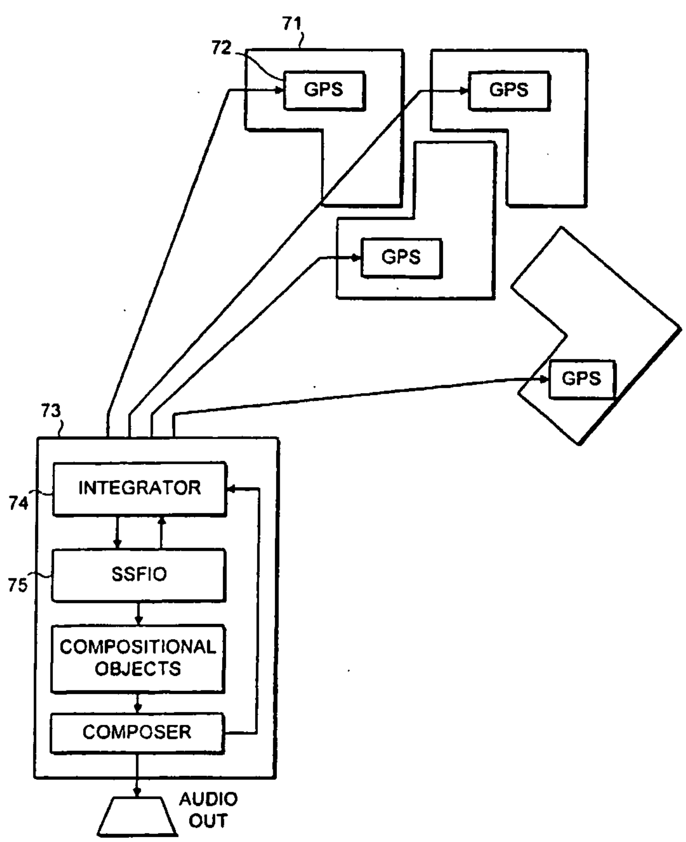

(57) Abstract: A generative sound system has a generative audio engine which is controlled or influenced by the absolute or relative location and/or orientation of a plurality of individual articles or units (71). The articles may, in a variety of embodiments, include collectable cards (86), building blocks (102), articles of furniture, ornaments and so on (128, 130, 132) or portable electronic devices such as mobile phones (60). Position and orientation information may be transmitted by means of a wireless message to the compositional engine, or the engine may determine that information in other ways such as by means of a proximity sensor (124, 136).

WO 01/86627 A2

(84) **Designated States** *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:
— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Automated Generation of Sound Sequences

This invention relates to methods and systems for automated generation of sound sequences, and especially (though not exclusively) of sound sequences in the form of music.

The automated creation of music has a long history, going back at least as far as Mozart's use of musical dice. One of the first musical works generated by a computer was L. Hiller's Illiac suite. Since that time, of course, the sophistication of computer-generated music or more generally audio sequences has increased substantially.

Systems for creating musical sequences by computer may conveniently be divided up into two areas, which have been called "non-generative" and "generative". Non-generative systems include deterministic systems which will produce the same sequences every time, along with systems that simply replay (perhaps in a random or other order) pre-composed sections of music. The vast majority of current systems which produce musical output make use of this type of approach, for example by selecting and playing a particular predefined sequence of notes at random when the key is pressed or a mouse button clicked. Generative Music Systems, on the other hand, may be considerably more complex. Such systems generate musical content, typically note by note, on the basis of a higher-level of musical knowledge. Such systems either explicitly or implicitly are aware of a variety of musical rules which are used to control or influence the generation of the music. In some systems, the rules may operate purely on the individual notes being generated, without imposing any form of higher order musical structure on the output; in

such systems, any musical order that arises will be of an emergent nature. More sophisticated systems may include higher-level rules which can influence the overall musical structure. Generative Music Systems will normally create musical content "on the fly", in other words the musical sequences are built up note by note and phrase by phrase, starting at the beginning and finishing at the end. This means that – in contrast with some of the non-generative systems - the musical content can be generated and played in real time: there is no need for example for the whole of the phrase to be generated before the first few notes of the phrase can be played.

For our present purposes, the essential features of a generative music system are that it generates musical content in a non-deterministic way, based upon a plurality of musical rules (which may either be implicit within the software or which may be explicitly specified by either the program writer or the user of the program). By analogy, a generative sound system produces non-deterministic sound sequences based upon sound-generation rules.

According to the present invention there is provided a generative audio system including a generative audio engine, the engine being controlled or influenced by the spatial positioning of a plurality of controlling items. It will be understood, of course, that the expression "spatial positioning" covers both absolute and relative positionings, as well as absolute and relative orientation.

The spatial positioning referred to may be defined by the relative spatial positioning between the individual controlling items, or any two or more of them. Alternatively, it may represent the relative spatial positioning between any one or more of the controlling items and a fixed or mobile base unit.

Where a base unit is supplied, the audio engine is preferable contained within or at least associated with the base unit. The base unit preferably also includes or is attached to an audio speaker.

In one embodiment, the base unit determines or is informed by the controlling units of each of the controlling units proximity to the base unit. The audio engine is then controlled or influenced on that basis.

Information on the spatial positioning of the controlling items may be transmitted to the audio engine either by a wire or a wireless connection. Alternatively, there may be no message transmission as such: instead, the spatial positionings may simply be detected by one or more suitable sensors.

The sounds or music could, further be controlled or influenced by the number, type, combination, proximity and composition of the controlling elements, and the ambient environment(s) in which they are located (e.g. light levels, humidity, temperature etc).

The invention extends to a game, and to a puzzle, incorporating a generative audio system as previously described. The individual controlling items are in one embodiment collectable items of some sort, such as cards, building blocks, small toys, items of jewellery or the like.

A device incorporating a generative music system of the invention may receive messages communicated, for example by wireless, relating to such matters as the musical activity, position or orientation of other devices. These other devices may be, for example, in the form of tags or tokens and the response of

the generative music system to the messages may be such as to indicate the relationship, in musical terms and/or positionally, of the tags or tokens to one another and/or the receiving device.

A method and system for automated generation of sound sequences, and applications of such method and system, according to the preferred embodiments of the invention will now be described, by way of example, with reference to the accompanying drawings, in which:

Figure 1 is a schematic representation of the system of the invention;

Figure 2 is illustrative of objects that are involved in a component of the system of Figure 1;

Figure 3 is a flow-chart showing process steps involved in control sequencing within the method and system of the invention;

Figure 4 is illustrative of operation of the method and system of the invention in relation to scale and harmony rules;

Figure 5 illustrates operation of the method and system of the invention in relation to the triggering of note sequences and their integration into a musical work as currently being composed and played;

Figure 6 shows in schematic form devices that each utilise the method and system of the present invention and are in wireless communication with one another and/or a base station;

Figure 7 is illustrative of an arrangement which includes a device that utilises the method and system of the present invention for providing sound output dependent on the position and/or orientation of other items.

Figure 8 is an embodiment in which sounds are generated in dependence upon the position of cards on a base unit;

Figure 9 shows schematically the cards used in the embodiment of Figure 8;

Figure 10 is another embodiment in which sounds are generated in dependence upon the position of stacking blocks on a base unit;

Figure 11 shows another embodiment in which cards can be slotted into a base unit;

Figure 12 shows another embodiment in which sounds are generated in dependence upon the position of objects within a room; and

Figure 13 shows yet a further embodiment in which the generation of music within a portable music player is influenced by jewellery worn by the user.

The method and system to be described are for automated generation of sound sequences and to integrate data presented or interpreted in a musical context for generating an output reflecting this integration. Operation is within the context of generation of musical works, audio, sounds and sound environments in real-

time. More especially, the method and system function in the manner of a `generative music system' operating in real-time to enable user-interaction to be incorporated into the composition on-the-fly. The overall construction of the system is shown in Figure 1 and will now be described.

Referring to Figure 1, the system involves four high-level layers, namely, an applications layer I comprising software components 1 to 5, a layer II formed by an application programmer's interface (API) 6 for interfacing with a music engine SKME that is manifest in objects or components 7 to 14 of a layer III, and a hardware device layer IV comprising hardware components 15 to 19 that interact with the music engine SKME of layer III. Information flow between the software and hardware components of layers I to IV is represented in Figure 1 by arrow-heads on dotted-line interconnections, whereas arrow-heads on solid lines indicate an act of creation; for example, information in the composed-notes buffer 11 is used by the conductor 12 which is created by the soundscape 8.

The applications layer I determines the look, feel and physical instantiation of the music engine SKME. Users can interact with the music engine SKME through web applications 1, or through desktop computer applications 2 such as those marketed by the Applicants under their Registered Trade Mark KOAN as KOAN PRO and KOAN X; the music engine SKME may itself be such as marketed by the Applicants under the Registered Trade Mark KOAN. Interaction with the engine SKME may also be through applications on other diverse platforms 3 such as, for example through mobile telephones or electronic toys. All applications 1 to 3 ultimately communicate with the music engine SKME via the API 6 which protects the internals of the music engine

SKME from the outside world and controls the way in which the applications can interact with it. Typically, the instructions sent to the API 6 from the applications 1 to 3 consist of commands that instruct the music engine SKME to carry out certain tasks, for example starting the composition and playback, and changing the settings of certain parameters (which may affect the way in which the music is composed/played). Depending on the needs of the individual applications, communication with the API 6 may be direct or via an intermediate API. In the present case communication to the API 6 is direct from the desktop computer applications 2, whereas it is via an intermediate browser plug-in API 4 and Java API 5 from applications 1 and 3 respectively.

The music engine SKME, which is held in memory within the system, comprises eight main components 7 to 14. Of these, SSFIO 7, which is for file input/output, holds a description of the parameters, rules and their settings used by algorithms within the engine, to compose. When the engine SKME is instructed via the API 6 to start composition/playback, a soundscape 8 is created in memory and this is responsible for creating a composer 10, conductor 12 and all the individual compositional objects 9 relating to the description of the piece as recorded in the SSFIO 7. The compositional objects are referred to by the composer 10 to decide what notes to compose next. The composed notes are stored in a number of buffers 11 along with a time-stamp which specifies when they should be played. The conductor 12 keeps time, by receiving accurate time information from a timer device 19 of level IV. When the current time exceeds the time-stamp of notes in the buffers 11, the relevant notes are removed from the buffers 11 and the information they contain (such as concerning pitch, amplitude, play time, the instrument to be used, etc.) is passed to the appropriate rendering objects 13. The rendering objects 13 determine

how to play this information, in particular whether via a MIDI output device 17, or as an audio sample via an audio-out device 18, or via a synthesiser engine 14 which generates complex wave-forms for audio output directly, adding effects as needed.

The hardware devices layer IV includes in addition to the devices 17 to 19, a file system 15 that stores complete descriptions of rules and parameters used for individual compose/playback sessions in the system; each of these descriptions is stored as an `SSfile', and many of these files may be stored by the file system 15. In addition, a MIDI in device 16 is included in layer IV to allow note and other musical-event information triggered by an external hardware object (such as a musical keyboard) to be passed into the music engine SKME and influence the composition in progress.

The system can be described as having essentially two operative states, one, a `dynamic' state, in which it is composing and the other, a `static' state, in which it is not composing. In the static state the system allows modification of the rules that are used by the algorithms to later compose and play music, and keeps a record encapsulated in the SSFIO component 7, of various objects that are pertinent to the description of how the system may compose musical works. The system is also operative in the dynamic state to keeps records of extra objects which hold information pertinent to the real-time composition and generation of these works. Many of these objects (the compositional objects 9 for example) are actual instantiations in memory of the descriptions contained in the SSFIO 7. Modification of the descriptions in the SSFIO 7 via the API layer II during the dynamic state, results in those modifications being passed down to the compositional objects 9 so that the real-time composition changes

accordingly.

Figure 2 shows a breakdown of the SSFIO component 7 into its constituent component objects which exist when the system is in its static and dynamic states; the system creates real-time versions of these objects when composing and playing. In this respect, SSfiles 20 stored each provide information as to `SSObject(s)' 21 representing the different *types* of object that can be present in the description of a work; these objects may, for example, relate to piece, voice, scale rule, harmony rule, rhythm rule. Each of these objects has a list of `SSFparameters' 22 that describe it; for example, they may relate to tempo, instrument and scale root. When an SSfile 20 is loaded into the music engine SKME, actual instances of these objects 21 and their parameters 22 are created giving rise to `SSFObjectInstance' 23 and `SSFParameterInstance' 24 as illustrated in Figure 2.

Referring again to Figure 1, the user interacts with the system through applications 1 to 3 utilising the services of the API 6. The API 6 allows a number of functions to be effected such as `start composing and playing', `change the rules used in the composition', `change the parameters that control how the piece is played' including the configuration of effects etc. One of the important advantages of the described method and system is the ability to trigger generative pattern sequences in response to external events. The triggering of a generative pattern sequence has a range of possible outcomes that are defined by the pattern sequence itself. In the event that a generative pattern sequence is already in operation when another trigger event is received, the currently operational sequence is ended and the new one scheduled to start at the nearest availability.

Generative pattern sequences allow a variety of musical seed phrases of any length to be used in a piece, around which the music engine SKME can compose in real time as illustrated in Figure 3. More particularly, the generative pattern sequence contains a collection of one or more note-control sub-patterns with or without one or more additional sequence-control sub-patterns. Three types of note-control sub-patterns can be created, namely: `rhythm' note-control sub-pattern containing note duration information, but not assigning specific frequencies to use for each note; `frequency and rhythm' note-control sub-pattern containing both note duration and some guidance to the generative music engine SKME as to the frequency to use for each note; and `forced frequency' note-control sub-pattern containing note duration, temporal positioning and explicit frequency information to use for each note. Sequence-control sub-patterns, on the other hand, can be used to specify the sequence in which the note-control sub-patterns are played, and each note-control sub-pattern may also specify ranges of velocities and other musical information to be used in playing each note. The music engine SKME allows the use of multiple sub-patterns in any generative pattern sequence.

Referring to Figure 3, the step 30 of triggering the generative pattern sequence acts through step 31 to determine whether there are any other sequence-control sub-patterns operative. If not, a note-control sub-pattern is chosen at random in step 32 from a defined set; each note-control sub-pattern of this set may be assigned a value that determines its relative probability of being chosen. Once it is determined in step 33 that the selected note-control sub-pattern is finished, another (or the same) note-control sub-pattern is selected similarly from the set. The generative pattern sequence continues to play in this manner until

instructed otherwise.

If the result of step 31 indicates that there is one or more sequence-control sub-patterns operative, then any sequence-control sub-pattern is chosen at random in step 34 from the defined set; each sequence-control sub-pattern may be assigned a value that determines its relative probability of being chosen. Once a sequence-control sub-pattern has been selected in step 34, it is consulted to determine in step 35 a sequence of one or more note-control sub-patterns to play. As each note-control sub-pattern comes to an end, step 36 prompts a decision in step 37 as to whether each and every specified note-control sub-pattern of the operative sequence has played for the appropriate number of times. If the answer is NO, then the next note-control sub-pattern is brought into operation through step 35, whereas if the answer is YES another, or the same, sequence-control sub-pattern is selected through repetition of step 34. As before, the generative pattern sequence continues to play in this manner until instructed otherwise.

Each sequence-control sub-pattern defines the note-control sub-pattern(s) to be selected in an ordered list, where each entry in the list is given a combination of: (a) a specific note-control sub-pattern to play, or a range of note-control sub-patterns from which the one to play is chosen according to a relative probability weighting; and (b) a value which defines the number of times to repeat the selected note-control sub-pattern, before the next sequence-control sub-pattern is selected. The number of repetitions may be defined as a fixed value (e.g. 1), as a range of values (e.g. repeat between 2 and 5 times), or as a special value indicating that the specified note-control sub-pattern should be repeated continuously.

Depending upon the note-control sub-pattern operational at any moment after a generative pattern sequence is triggered, various rules internal to the music engine SKME may be used to determine the exact pitch, duration and temporal position of the notes to be played. For example, if a `rhythm' note-control sub-pattern is in operation at a particular point in the generative pattern sequence, then the scale rule, harmony rule and next-note rule within the music engine SKME for that `triggered voice' will be consulted to obtain the exact notes. Alternatively, if the `forced frequency' note-control sub-pattern is operational, no internal rules need be consulted since all the note information is already specified. Furthermore, for the case of `frequency and rhythm', the music engine SKME combines the given frequency offset information with its rules and other critical information such as the root of the current scale and range of available pitch values for the voice in question.

The rules and other parameters affecting composition (e.g. tempo) within the music engine SKME are defined in memory, specifically within the SSFIO 7, and its real-time instantiation of the compositional objects 9. Use of rules and parameters within the music engine SKME form part of the continual compositional process for other voice objects within the system. Figure 4 illustrates this more general process based on examples of scale and harmony rules shown at (1) and (2) respectively.

Referring to Figure 4, the scale rule is illustrated at (1) with shaded blocks indicating a non-zero probability of choosing that interval offset from a designated scale root note. The larger the shaded block, the greater the probability of the system choosing that offset. Thus, for this example, the

octave Ove, major third M3 and fifth 5 are the most likely choices, followed by M2, 4, M6 and M7; the rest will never be chosen. Sequences that may be generated by the system from this are shown below the blocks, and in this respect the octave has been chosen most often followed by the major third and the fifth. With the scale root set in the system as C, the resulting sequence of notes output from the system in this example are C,E,C,D,G,A,E,D,C,G,E,B,C,F, as illustrated at (1) of Figure 4.

The harmony rule defines how the system may choose the pitches of notes when other notes are playing, that is to say, how those pitches should harmonise together. In the example illustrated at (2) of Figure 4, only the octave and major second are indicated (by shading) to be selected. This means that when the pitch for a voice is chosen, it must be either the same pitch as, or a major second from, all other notes currently being played.

For the purpose of further explanation, consideration will be given to the example represented at (3) of Figure 4 involving three voice objects V1-V3. The rhythm rules applicable to the voice objects V1-V3 in this example, give rise to a generated sequence of notes as follows: voice V1 starts playing a note, then voice V2 starts playing a note, then voice V3 starts playing a note, and then after all notes have ended, voice V2 starts playing another note, followed by voice V1 and then voice V3. With this scenario, the note from voice V2 must harmonise with that of voice V1 and the voice V3 note must harmonise with that of voice V2. If in these circumstances the voice V1 is, as illustrated by bold hatching, chosen with a pitch offset of a fifth from the scale root, the pitch for voice V2 must either be the same as (Ove) or a major second above (M2) the fifth. In the case illustrated, it is chosen to be the same, and so the

fifth is chosen too. When voice V3 starts playing it must harmonise with both voices V1 and V2, so the pitch chosen must be the same as, or a major second above that of voices V1 and V2. As illustrated, the system chooses voice V3 to be a major second above, therefore giving pitch offset M6 from the scale root.

After voice V3 all notes end, and the next note begins, as illustrated at (4) of Figure 4 with voice V2. This next note by voice V2 is governed by the next-note rule used by voice V2, and the last note played by voice V2. According to this rule, the system chooses pitch offset M2 for voice V2, and then harmonises voices V3 and V1 with it by choice of a major second for both of them. With the scale root set in the system to C, the entire generated sequence accordingly follows that indicated at (5) of Figure 4, where `S' denotes a note starting and `E' a note ending.

Thus, when sequences are generated in response to an external trigger, the actual pitches and harmonisation of that sequence is determined by the composer 10 using several items of information, namely: (a) the note-control sub-pattern operational at that moment; (b) the scale, rhythm, harmony and next-note rules depending upon the type of the note-control subsequence; and (c) any piece-level rules which take into account the behaviour of other voices within the piece.

When the music engine SKME is in dynamic (i.e. composing and playing) mode, it typically contains a number of voice compositional objects 9. The composer 12 composes a sequence of notes for each of these and makes sure they obey the various rules. The process involved is illustrated in the flow diagram of Figure 5.

Referring to Figure 5, the music engine ~~SKME~~ responds to an external trigger applied at step 51, and the API 6 through step 52 instructs a voice 1 in step 53 to register that it must start a sequence. Voice 1 and the voices 2 to N in step 54, have their own rules, and the composer 10 ensures that the relevant rules are obeyed when utilising any of the voices 1 to N. More particularly, the composer 10 responds in step 55 to the instruction of step 53 for voice 1 to start a sequence, by starting the generative pattern sequence sub-system of Figure 3. This sends note-control sub-sequences to the trigger voice (voice 1 in this example), but the composer 10 makes sure the resulting notes harmonise with the other voices in the piece. The outcome via the conductor 12 in step 56 is played in step 57.

The generative pattern sequence triggered will play forever, or until the system is instructed otherwise. If a sequence control sub-pattern is used to define a generative pattern sequence such that the final note control sub-pattern is one which plays silence (rest notes) in an infinite loop, then when this pattern sequence is selected, the voice will become effectively `inactive' until another trigger is detected. Further triggering events for the same generative pattern sequence may sound different as the process is generative, or since the rules in use by the piece or the scale of the trigger voice, its harmony or next note rules may have changed (either via interaction through the API 6 or via internal music engine SKME changes).

The sounds used to `render' each note, whether from triggered sequences or generative voices may be played either through the MIDI sounds or the samples of the rendering objects 13, or via software of the synthesiser engine 14 which

may add digital signal processing effects such as, for example, filter sweeps and reverberation and chorus. The entire process can be used to generate musical event information that is then fed into, and may thus control, other processing units within the system such as synthesiser related units allowing the triggering of generative sound effects. Voices can also be added which make use of the software synthesiser engine 14 to generate non note-based effects such as sound washes and ambient environmental sounds, such as chimes, wind and other organic sounds.

The method and system of the invention is applicable with advantage in a networked system. In particular, it may be used for the purpose of networked musical "jamming" (joint composition). Figure 6, illustrates an example of their application in this context.

Referring to Figure 6, two wireless networked devices 60 (for example mobile phones, portable music systems, gaming consoles etc) each include a generative music system (not shown in full detail) of the form described above, the devices 60 being in this respect typical of a multiplicity of digital devices linked together for wireless communication in the relevant network. Each device 60 is in wireless communication with an (optional) base unit 65. Each device 60 also includes an integrator 61 that receives information by reception of messages transmitted to it by wireless from the base unit 65 (and other network devices), and also internally by information from within its own generative music system. Each mobile device 60 has some knowledge of its physical location, either its absolute location or alternatively its relative location, relative to other mobile units 60 or to the base unit 65. Absolute location could be determined either by a GPS (Global Positioning System) within each individual

device, or alternatively by means of messages transmitted from the central unit. In a cellular system, for example, approximate knowledge of the cell in which each device 60 is operating will typically be known to the system because the device will be connecting into the fixed network infrastructure by means of a particular access point; in such a case, of course, the base units 65 will in practice be a plurality of fixed access nodes. Alternatively, the units 60 may obtain some knowledge of their absolute position according to the frequency or some other characteristic of the communications with the central unit 65.

The musical sounds generated by the devices 60 may also be controlled or influenced by the relative positions of the devices 60 themselves and/or the central unit 65. The information on relative positions could be achieved by message passing either directly between the devices 60 or between each device 60 and the central unit 65. Alternatively, some other method of proximity detection could be used: for example, each device 60 could be sensitive to the amount or frequency of the electromagnetic radiation emitted from other similar devices or from the base unit. A high level of radiation detected (or alternatively a strong signal) may imply proximity, with weaker signals implying lack of proximity. By making use of detected signal strengths, the devices 60 are able, in one embodiment, to determine their position or relative position without the need for message exchanges.

Each device 60 receives information from the SSFIO 62 and also from the composer 64 that is linked to it through the compositional objects 63, of its generative music system. The information from the SSFIO 62 describes its current musical `behaviour', whereas that from the composer 64 describes the current state of the musical output from the device 60. The information from all

three sources is used within the device 60 to make changes to the SSFIO 62 so as to affect the future musical behaviour of the device 60.

Since each device 60 has its own generative music engine, any wireless messages passing between the devices 60, or between a device 60 and the central unit 65, requires only a very small bandwidth. Such messages may be effectively small files which can: define explicitly the compositional rules or other elements to be used to compose/generate audio in the relevant receiving device 60; (b) describe instructions which are effective to modify the compositional rules in an integrative fashion within the receiving device 60; and (c) effect the changes required in near real-time in the receiving device 60. The messages may moreover contain small audio sample files or descriptions of sound processing units or effects (e.g. synthesiser unit descriptors).

Turning now to Figure 7, this shows an alternative embodiment involving the use of units that communicate positional and/or orientational information to a generative music system, so that the output of the system is dependent upon that information.

Referring to Figure 7, a plurality of tokens or units 71 each containing a positioning system e.g. a Global Positioning System (GPS) 72 transmit wireless messages at intervals to a device 73 that corresponds to the device 60 of the arrangement of Figure 6. More particularly, the device 73 incorporates an integrator 74 that operates in conjunction with the generative music system of the device 73 to monitor the incoming messages from the individual units or tokens 71; the units 71 may be interrogated in turn to prompt the transmission of the messages. The messages include positional and/or orientational information concerning the individually identified units 71 and the integrator 74

passes appropriately modified messages to the SSFIO 75 of the generative music system. The result is that the musical output of the device 73 is dependent upon the positions and/or orientations of the various tokens or units 71, in such a way that arranging the units 71 in different geometric combinations and alignments achieves different audio effects or compositions.

Alternatively, the Positioning System feature of the units 71 may be omitted, and each unit may instead be arranged to determine its absolute or relative position and/or orientation in some other way. For example, positional/inclination sensors may be provided on each of the units 71, these being used to supply the information that is needed to control or influence sound/music generation within the device 73. It would also be possible for the units 71 to determine their own relative positions and inclinations – for example by means of proximity sensors – with the resultant information being reported back by means of a message or messages sent to the device 73.

As with Figure 6, in some embodiments of Figure 7, message passing between the units 71 and the device 73 may not be necessary. Instead, the device 73 may have some other way of determining the absolute or relative position and/or orientation of the units 71. The device 73 may for example itself include proximity sensors, enabling it to sense the position and/or distance of each of the units 71.

In a local environment, the approach described generally above with reference to Figure 7 may be used in a number of specific applications, some of which will be described in more detail below with reference to Figures 8 to 13.

In Figure 8, the music-generating device 73 is contained within a flat base unit 80 which is arranged to be connected to an external loudspeaker 82 by a lead 84. Alternatively, the speaker could be built into the base unit 80 itself.

The flat upper part of the base unit is arranged to receive cards or tokens 86, either positioned anywhere the user desires on the surface or placed into pre-formed bays or slots (not shown). As shown in Figure 9, each card 86 includes, preferably on its rear surface, a card locator 90 which, in association with suitable electronics (not shown) within the base unit, enables the system to be able to determine exactly where on the surface each individual card has been located. The card locator 90 could be as simple as a metal loop to be read by the base unit, or a magnetic or metallic strip, or a smart card.

The generative musical sound engine within the base unit 80 is controlled or influenced by the number of cards on the surface, the card types, and the cards absolute or relative positions and/or orientations.

Each type of card may control or influence the musical sound being generated in its own individual way. One card might, for example, control the bass voice, another the guitar, another the drums and so on. Other types of card may influence the musical sound generation in other ways, for example controlling the volume of one or more voices, pitch, timbre, rhythm and so on. Individual cards may also contain pre-defined musical or sound "templates" to produce a basic or skeleton composition which may then be influenced or augmented by other cards.

One card type 98 may be a master card, defining the overall rule set. This may

need to be placed in a specific region of the based unit's upper surface, for example the left hand corner, as shown in Figure 8. In addition to including templates, where appropriate, the master card 98 could also define or control the contribution to the composition of the other cards. It might, for example, instruct the system that cards of type 1 should be treated as the bass line, cards of type 2 as the guitar line and so on. In that way, the use of different master cards may fundamentally affect the overall composition.

The cards 86 and/or 98 may be user-programmable, provided of course that the card location element 90 includes some type of non-volatile memory. Users could for example download template definitions, master card definitions, sample sound or music descriptions and musical rules, individual parameter values and so on from a central server (e.g. via the Internet).

The front surface of each card may include a printed design and/or text which differs according to card type. The cards may then be collected and/or swapped.

It will be understood of course that the embodiment of Figure 8 may be extended to other toy or collecting scenarios, in which the cards 86,98 may be replaced with books, toys or other collectable units.

Some further embodiments will now be described with reference to Figures 10 to 13. It will be understood that each individual movable element within those embodiments may have any of the features, characteristics or functionality of the cards 86,98 of Figure 8 or the units 71 of Figure 7.

Turning first to Figure 10, there is shown an alternative embodiment in which the cards 86,98 are replaced with building blocks 102. The blocks may be connectable in some way – for example in the manner of Lego™ blocks or Sticklebricks™ – to enable towers 104 to be built. The blocks are placed on, or may be securable to, a base unit 100 which operates in a manner similar to that of the base unit 80 of Figure 8.

The musical sounds generated by the system may in this embodiment depend on the position of bricks in the third dimension, as well as the positioning on the surface of the base unit 100. The height of the towers 104, their locations, rotational positions, angles of attachment the type of blocks within them may all affect the sound or music generation.

In this and other embodiments, the relative positions of the blocks/units may determine the flow of the music – for example each block may represent a single bar of music with the bars being played according to a left-to-right or top-to-bottom progression.

In this and other embodiments the blocks/units may include sensors (not shown) for sensing one or more details of the ambient environment (e.g. light levels, humidity, temperature etc). The engine may further be controlled or influenced by the outputs of one or more such sensors. For example, the engine could (at least partially) control the composition in dependence upon changing light patterns falling into a pile, or tower, of bricks.

The device shown in Figure 10 may be configured so that the audio output reflects how close the users is to a desired configuration (for example in solving

a puzzle associated with the blocks). A similar concept may be used in conjunction with the embodiment of Figure 8: for example, the output may depend upon how close the user has got to laying out the cards to a predefined configuration, initially unknown to the user. The positioning of blocks may be an interesting way to control the temporal evolution of the piece of music. In such a way the user can mix and improvise by repositioning blocks at particular points in time.

In an alternative embodiment (not shown) the individual objects could take the form of balls, preferably within an enclosed container such as a sphere. By jiggling the balls about within the container, the user could create a kaleidoscope of sound as the positions of the balls continually vary.

Turning next to Figure 11, a base unit 110 has a plurality of parallel slots 112 for receiving individual cards such as the cards 86,98 of Figure 8. The base unit is connected to an external loudspeaker 111; alternatively, the loudspeaker may be integrated within the unit.

A collector of the cards 86,98 places them, as desired, into the various slots 112 in order to control or influence the musical sounds being played by the speaker 111. The system detects how many cards have been placed into slots, the card types, and the card locations in order to control the music or sound generation.

In another embodiment, shown in Figure 12, a generative sound or music system is contained within a speaker enclosure 120 having a speaker 112. The enclosure has a sensor 124 enabling it to sense or to receive signals from movable articles of furniture, crockery, ornaments and the like. In the example

shown, the sensor receives wireless signals 126 from a vase 128, a chair 130 and coffee cups 132. As in the previous embodiments, the sound or music generator is controlled or influenced by the absolute or relative positioning within the room of the various movable items.

Instead of receiving signals from each of the items, the sensor 124 could detect their location or proximity in some other way. Where message-passing is involved, provision may be made for the individual items to pass messages between themselves, for example as indicated by the dotted line 134, enabling more complex interactions to take place. If, for example, the coffee cups on the table can between them determine their relative positions and/or locations quite precisely, that information may then be passed back to the sound or music generator to influence the overall composition, without any need for the sensor 124 to distinguish, on its own, between the locations of the two cups.

The enclosure 120 may be static within the room, and may for example form part of a hi-fi system. The system could then generate and play background music and/or sounds in dependence upon the configuration of the furniture and/or other items within the room.

Alternatively, the enclosure 120 could be portable, and could be carried on the owner's person. Then, as the owner walks around the room, the music and/or sounds generated by the system will automatically vary.

In another alternative, the enclosure 120 could be dispensed with, and one or more speakers built into the individual objects themselves.

Figure 13 shows yet a further embodiment in which the generative engine is contained within a small belt-mounted unit 130 which is connected by means of a lead 132 to a pair of earphones 134. A sensor 136 on the unit 130 detects the presence and/or location of jewellery/clothing or other items being worn by the user, for example a watch 138 and a bracelet 140. Each piece of jewellery includes a transponder 142, 144 permitting communication with the detector 136, as indicated by the dotted line 146. As with the previous embodiments, messages may be passed or alternatively the sensor 136 may simply detect the presence/absence and the proximity of individual items of jewellery/clothing in some other way. The items of jewellery/clothing might be on another person fitted with a similar system, so people can influence the audio experienced by the other simply by moving themselves (and their belongings) in relation to that person.

For more complex effects, message-passing between the individual items may occur, as indicated by the dotted line 148, as well as between the items and the sensor 136. Messages transmitted to the sensor could control the generative engine based upon, for example, the number, type and sequence of individual beads or other elements on the bracelet 144. It will be understood of course that the sounds or music could more generally be controlled or influenced by the number, type, relative locations, orientation, proximity and composition of the various controlling elements.

CLAIMS:

1. A generative audio system including a generative audio engine, the engine being controlled or influenced by the spatial positioning of a plurality of controlling items.

2. A generative audio system as claimed in claim 1 in which the engine is controlled or influenced by the relative positionings, between themselves, of the controlling items.

3. A generative audio system as claimed in any one of the preceding claims in which the engine is controlled or influenced by the absolute or relative orientations of the controlling units.

4. A generative audio system as claimed in any one of the preceding claims in which the engine is controlled or influenced by the proximity, one to another, of the controlling units.

5. A generative audio system as claimed in any one of the preceding claims including a base unit, the engine being controlled or influenced by the relative positionings of the controlling items with respect to the base unit.

6. A generative audio system as claimed in claim 5 in which the audio engine is located within the base unit.

7. A generative audio system as claimed in any one of claims 1 to 6 in which at least some of the controlling items include their own audio engines.

8. A generative audio system as claimed in claim 5 or claim 6 in which the controlling items inform the base unit of their spatial positionings.

9. A generative audio system as claimed in claim 5 or claim 6 in which the base unit includes a sensor which detects the spatial positionings of the controlling items.

10. A generative audio system as claimed in claim 5 or claim 6 in which the controlling items determine their own relative positionings, and advise the base unit accordingly.

11. A generative audio system as claimed in claim 5 or claim 6 in which the engine is controlled or influenced by moving the base unit relative to stationary controlling items.

12. A generative audio system as claimed in claim 5 or claim 6 in which the engine is controlled or influenced by moving the controlling items relative to a stationary base unit.

13. A generative audio system as claimed in any one of claims 1 to 12 in which the controlling items are cards.

14. A generative audio system as claimed in any one of claims 1 to 12 in which the controlling items are building blocks.

15. A generative audio system as claimed in any one of claims 1 to 12 in

which the controlling items are mobile phones.

16. A generative audio system as claimed in any one of claims 1 to 12 in which the controlling items are pieces of jewellery.

17. A generative audio system as claimed in any one of claims 1 to 12 in which the controlling items are items of clothing.

18. A generative audio system as claimed in any one of claims 1 to 12 in which the controlling items are items of furniture, ornaments or other household items.

19. A generative audio system as claimed in any one of claims 1 to 12 in which the controlling items are personal, portable music systems.

20. A generative audio system as claimed in claim 13 or claim 14 when dependent upon claim 5, in which the engine is controlled or influenced by the spatial positioning of controlling items placed on the base unit.

21. A generative audio system as claimed in claim 13 including a slotted base unit into which the cards are received.

22. A generative audio system as claimed in any one of the preceding claims in which the controlling items are of a plurality of differing types, each type controlling or influencing the engine in a different way.

23. A generative audio system as claimed in claim 22 including a master

controlling item which controls how the engine responds to controlling items of a given type.

24. A generative audio system as claimed in any one of the preceding claims in which the controlling items include user-changeable control information defining the manner in which they control or influence the audio engine.

25. A game incorporating a generative audio system as claimed in any one of the preceding claims.

26. A puzzle incorporating a generative audio system as claimed in any one of the preceding claims.

27. A generative audio system as claimed in any one of claims 1 to 24 in which at least some of the controlling items include sensors, the engine being further controlled or influenced by the sensor outputs.

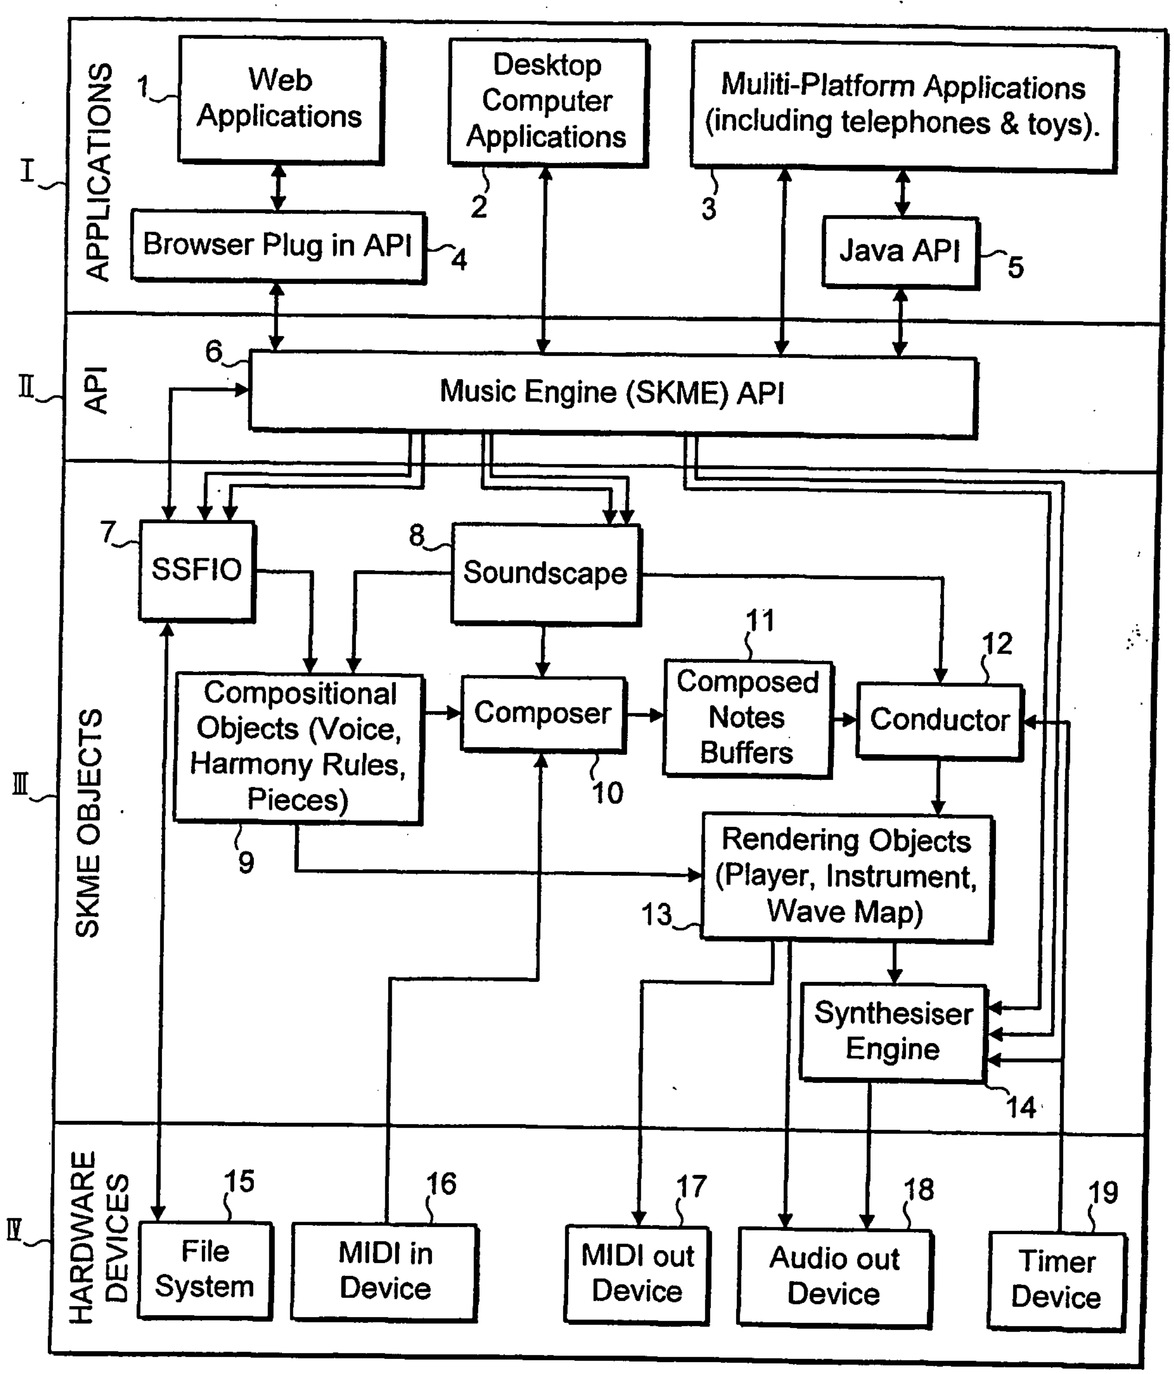
APPLICATIONS
I
1
Web Applications
Desktop Computer Applications
2
Muliti-Platform Applications (including telephones & toys).
3
Browser Plug in API
4
Java API
5
API
II
6
Music Engine (SKME) API
SKME OBJECTS
III
7
SSFIO
8
Soundscape
11
12
Compositional Objects (Voice, Harmony Rules, Pieces)
9
Composer
10
Composed Notes Buffers
Conductor
Rendering Objects (Player, Instrument, Wave Map)
13
Synthesiser Engine
14
HARDWARE DEVICES
IV
15
File System
16
MIDI in Device
17
MIDI out Device
18
Audio out Device
19
Timer Device

FIG. 1

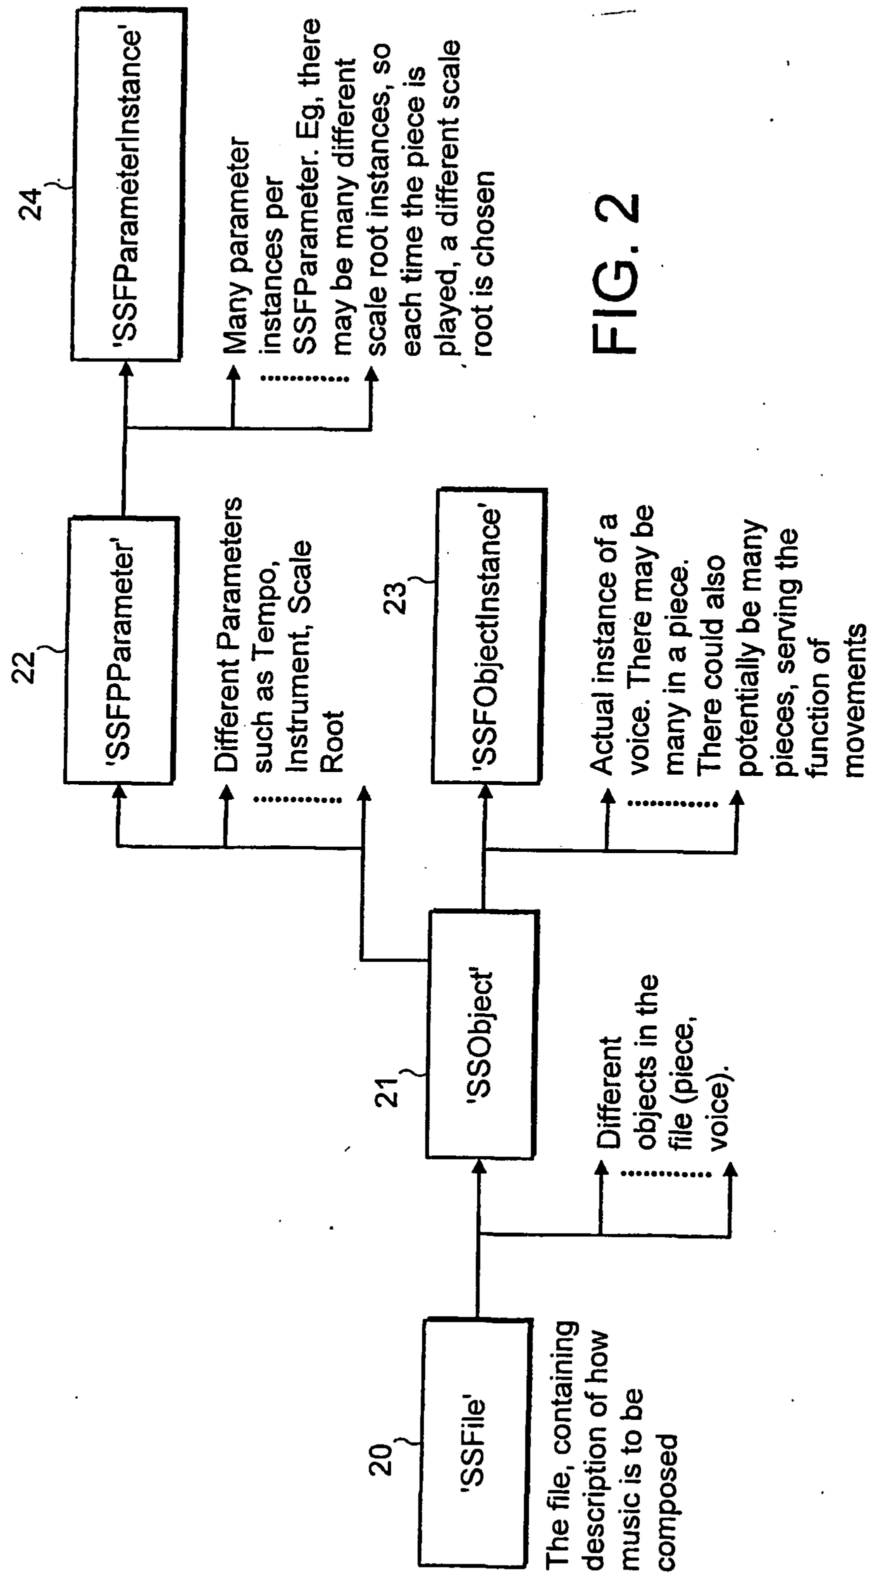
20
'SSFile'
The file, containing description of how music is to be composed
21
'SSObject'
Different objects in the file (piece, voice).
22
'SSFPParameter'
Different Parameters such as Tempo, Instrument, Scale Root
23
'SSFObjectInstance'
Actual instance of a voice. There may be many in a piece. There could also potentially be many pieces, serving the function of movements
24
'SSFParameterInstance'
Many parameter instances per SSFParameter. Eg, there may be many different scale root instances, so each time the piece is played, a different scale root is chosen

FIG. 2

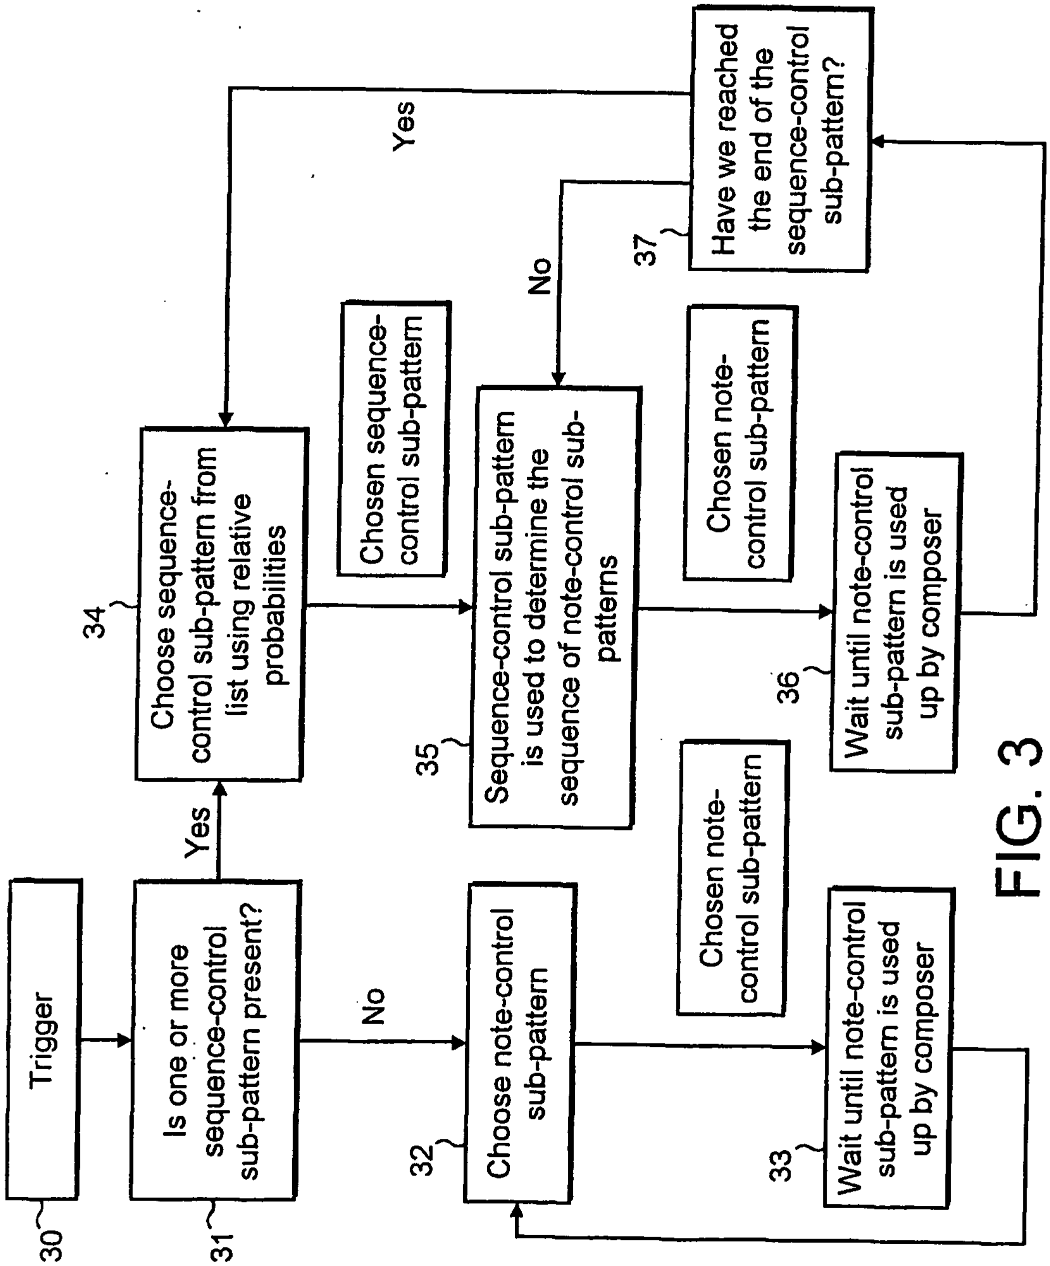
Trigger
30
Is one or more sequence-control sub-pattern present?
31
Yes
No
Choose sequence-control sub-pattern from list using relative probabilities
34
Chosen sequence-control sub-pattern
Sequence-control sub-pattern is used to determine the sequence of note-control sub-patterns
35
Chosen note-control sub-pattern
Wait until note-control sub-pattern is used up by composer
36
Have we reached the end of the sequence-control sub-pattern?
37
Yes
No
Choose note-control sub-pattern
32
Chosen note-control sub-pattern
Wait until note-control sub-pattern is used up by composer
33

FIG. 3

1 An example Scale Rule:

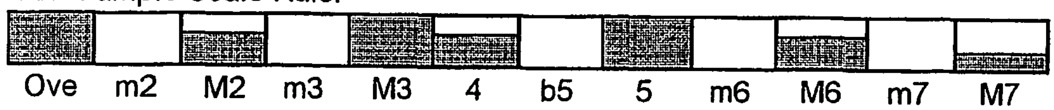

Relative to a scale root, a sequence of notes chosen from this scale rule might be:
Ove, M3, Ove, M2, 5, M6, M3, M2, Ove, 5, M3, M7, Ove, 4
With the scale root set at C, this sequence becomes:
C, E, C, D, G, A, E, D, C, G, E, B, C, F

2 An example Harmony Rule:

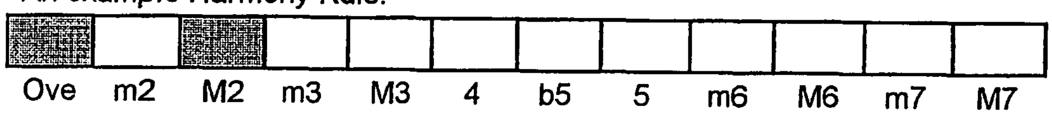

3

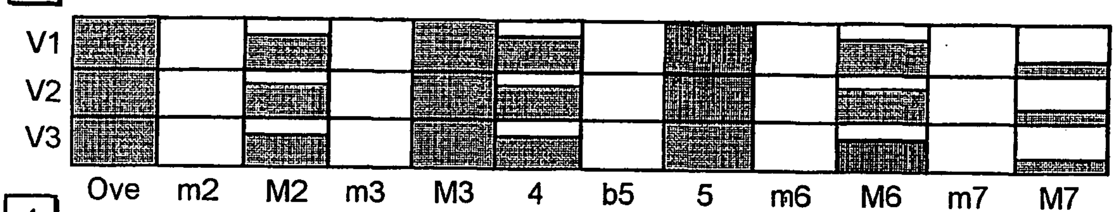

4

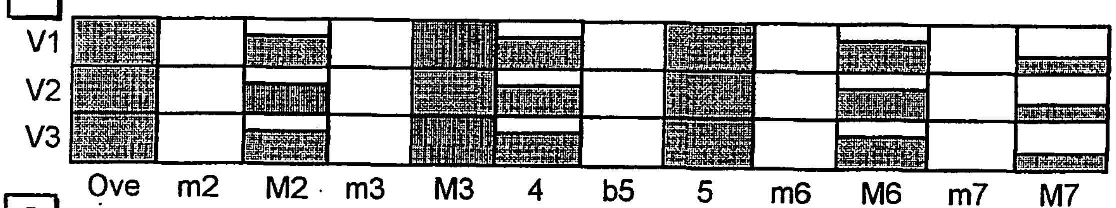

5

Time Sequence:
S:V1:G, S:V2:G, S:V3:A,
E:V1. E:V2, E:V3
S:V2:D, S:V3:E, S:V1:E
E:V2, E:V3, E:V1

FIG. 4

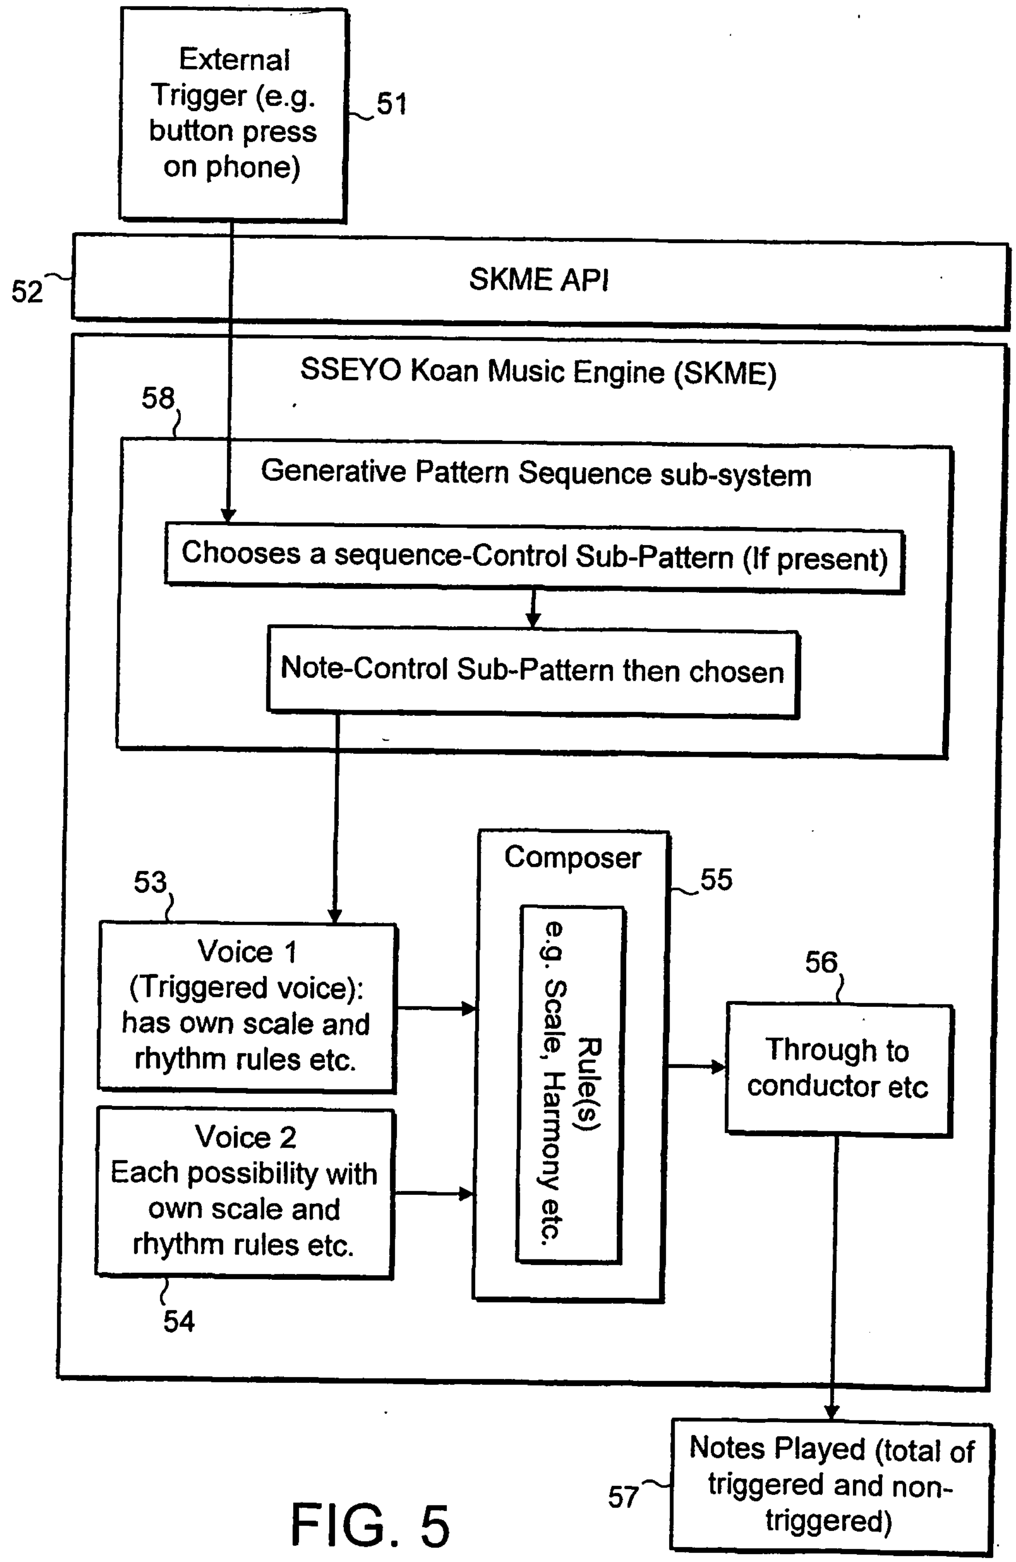

External Trigger (e.g. button press on phone)
51
52
SKME API
SSEYO Koan Music Engine (SKME)
58
Generative Pattern Sequence sub-system
Chooses a sequence-Control Sub-Pattern (If present)
Note-Control Sub-Pattern then chosen
53
Voice 1 (Triggered voice): has own scale and rhythm rules etc.
Voice 2 Each possibility with own scale and rhythm rules etc.
54
Composer
55
Rule(s) e.g. Scale, Harmony etc.
56
Through to conductor etc
57
Notes Played (total of triggered and non-triggered)

FIG. 5

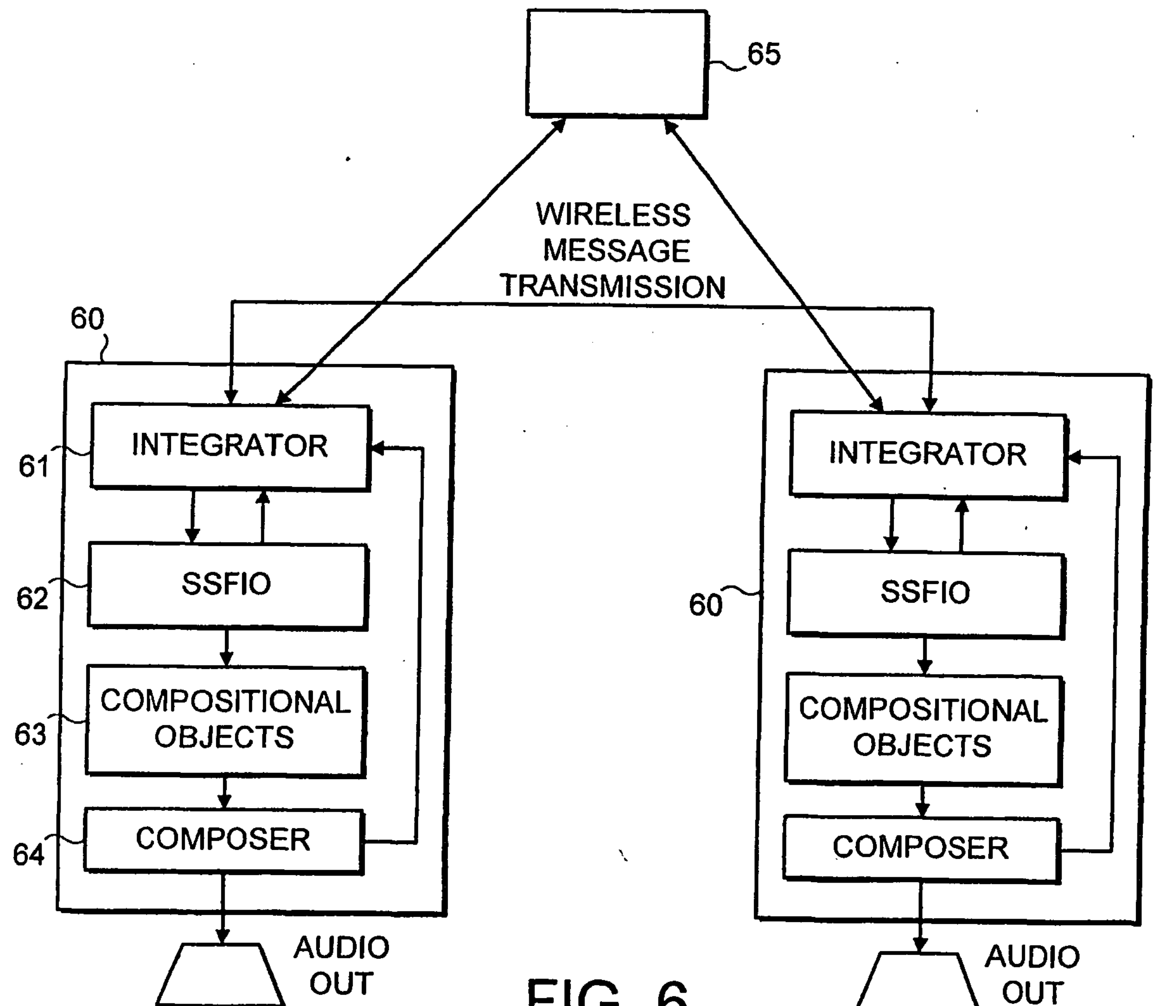
65
WIRELESS
MESSAGE
TRANSMISSION
60
61
INTEGRATOR
62
SSFIO
63
COMPOSITIONAL
OBJECTS
64
COMPOSER
AUDIO
OUT
60
INTEGRATOR
SSFIO
COMPOSITIONAL
OBJECTS
COMPOSER
AUDIO
OUT

FIG. 6

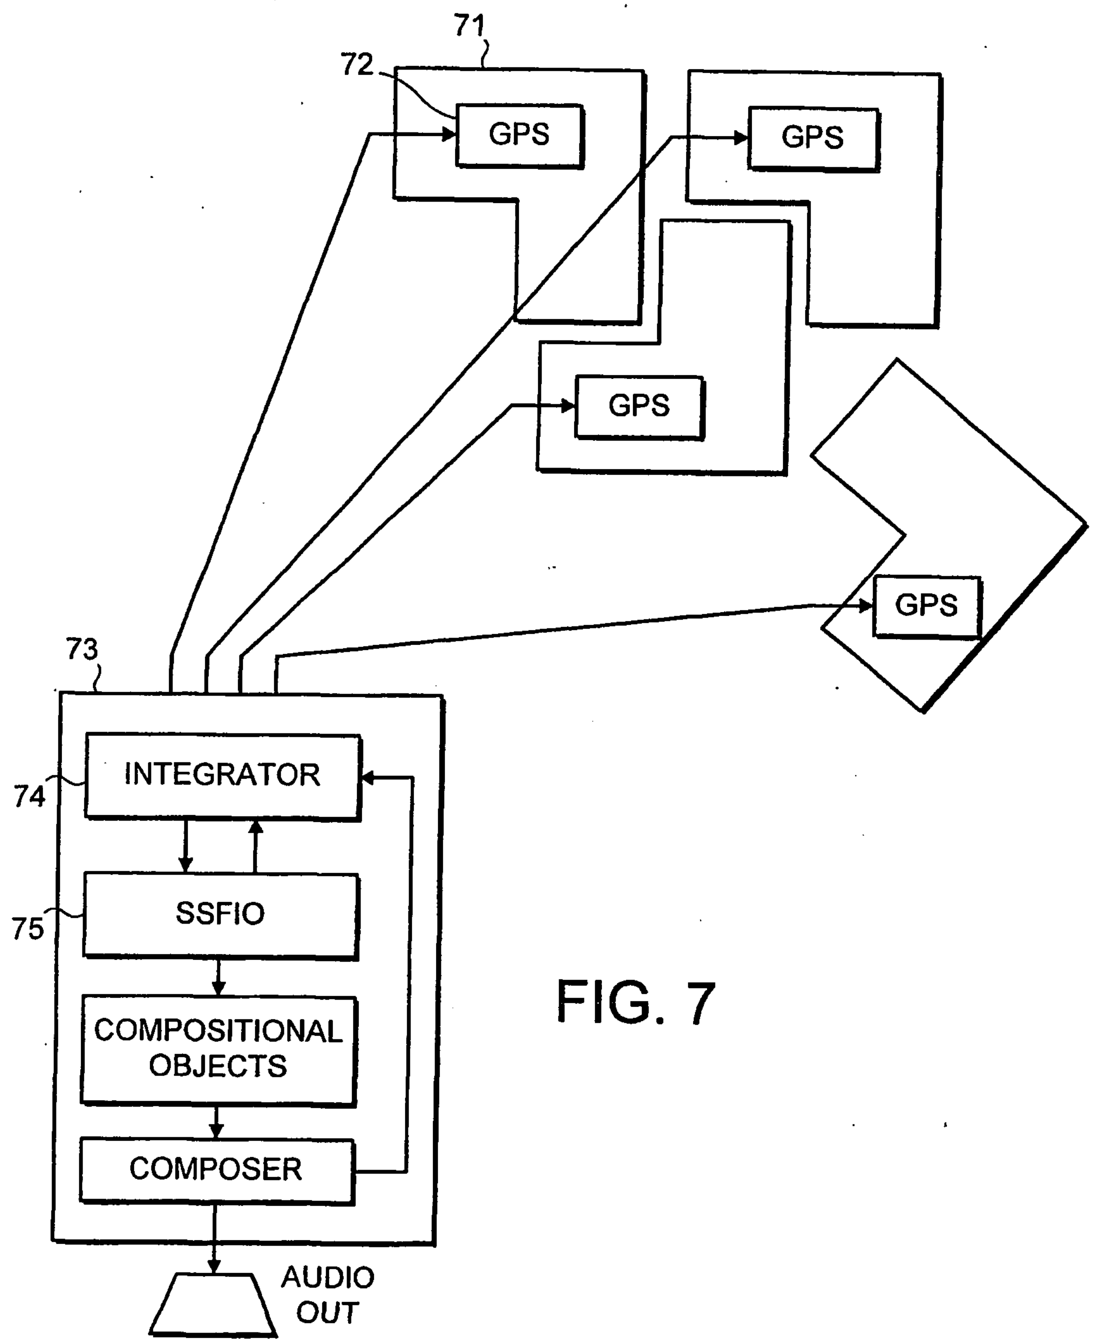
71
72
GPS
GPS
GPS
GPS
73
INTEGRATOR
74
SSFIO
75
COMPOSITIONAL OBJECTS
COMPOSER
AUDIO OUT

FIG. 7

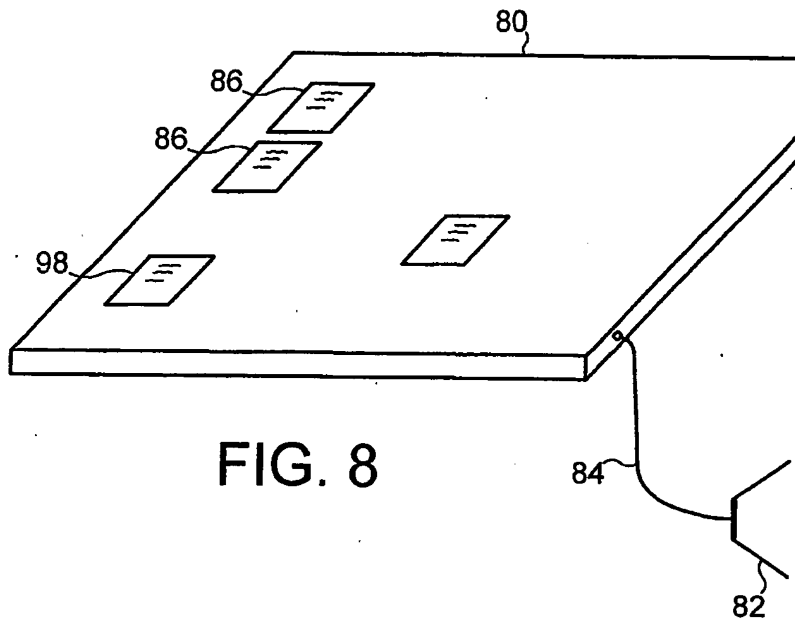

FIG. 8

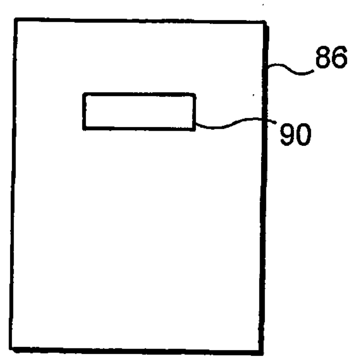

FIG. 9

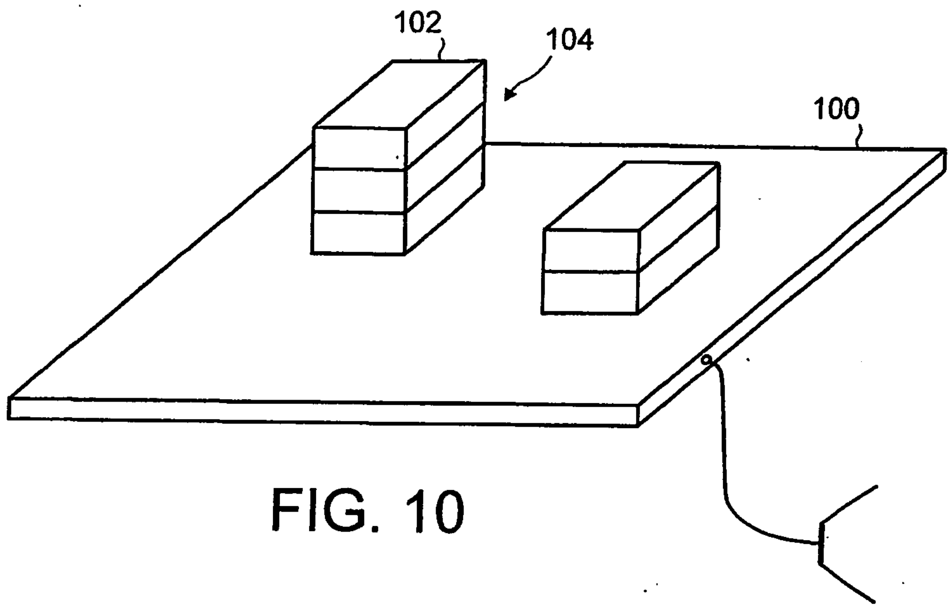

FIG. 10

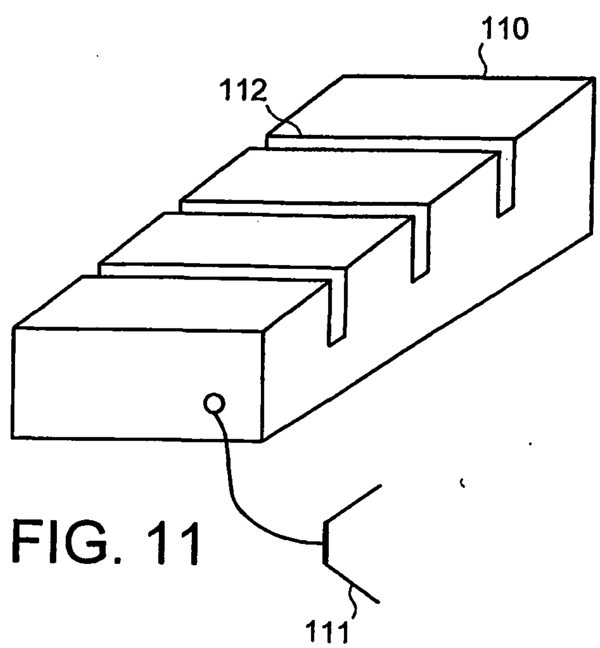

FIG. 11

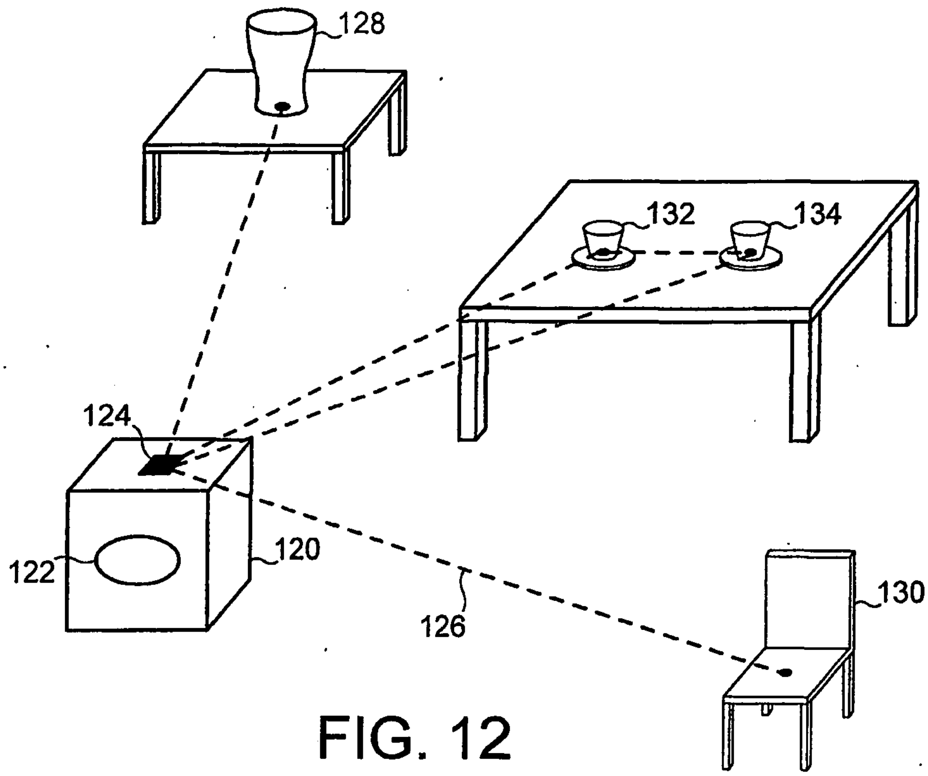
128
132
134
124
120
122
126
130

FIG. 12

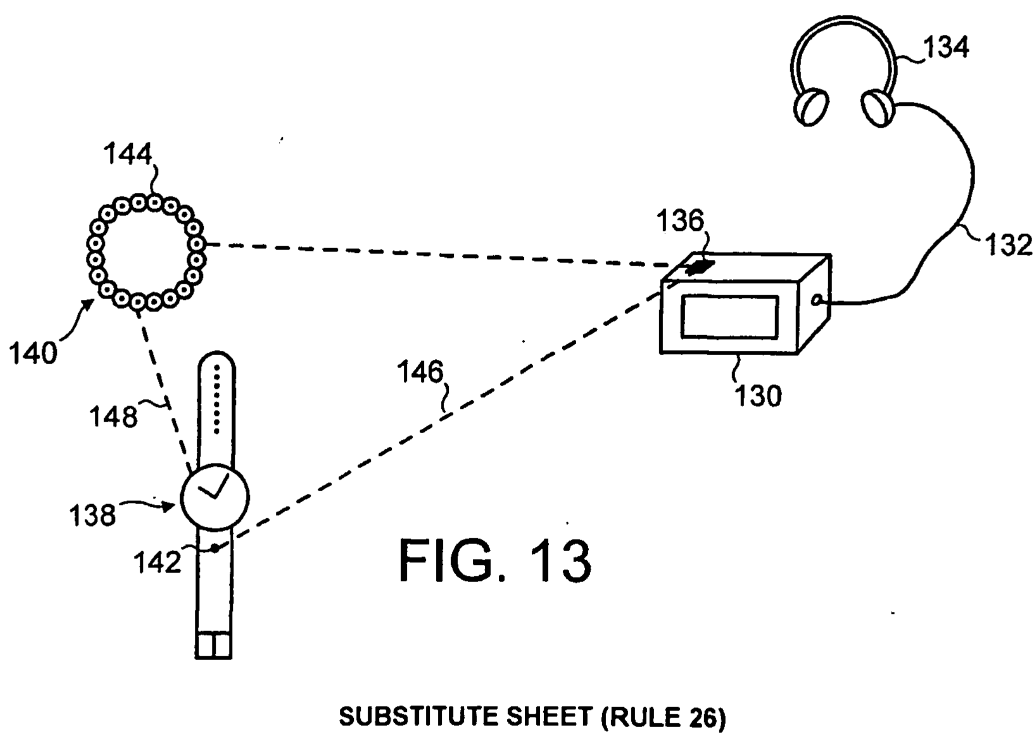
134
144
136
132
140
146
130
148
138
142

FIG. 13